JPL D-12771

Correlation of the Magellan Flight PFR History With Ground-Test Results

Prepared by:
Charles C. Gonzalez

Approved:

John E. Koch, Task Manager
Flight Performance Assessment RTOP

January, 1996

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

PREFACE

The NASA Unmanned Flight Anomaly Reports present the results of a series of analyses of in-flight hardware anomalies which have occurred on the Jet Propulsion Laboratory (JPL), Goddard Space Flight Center (GSFC), and Air Force unmanned space programs. All of these analyses are funded by NASA's Office of Safety and Mission Assurance (Code QT) under Research Technology Operation Plan (RTOP) 623-63-03, entitled *Flight Anomaly Characterization* (FAC). The objective of these analyses is to search for meaningful characterizations of in-flight anomaly data relating to trends, patterns, or similarities that can be exploited to improve Product Assurance Program processes, ultimately leading to reduced numbers of anomalies on future unmanned flight programs.

For further information on the content of this report, contact Charles Gonzalez at (818) 354-2500. For additional copies of this document, contact the JPL Document Vellum Files at extension 4-6222.

ABSTRACT

Correlation of the Magellan Flight PFR History with Ground-Test Results

This NASA Unmanned Flight Anomaly Report describes the results of an investigation and characterization of in-flight and ground-test anomalies for the JPL Magellan Program. The primary objective of this characterization and analysis process is to find groups of anomalies that share some common characteristics and that might be eliminated on future space flight programs by an improved product assurance process.

ACKNOWLEDGMENT

The author wishes to thank Art Brown, manager of the entitled *Flight Anomaly Characterization* (FAC) task for his many helpful suggestions on the writing of this report. In addition, David Oberhettinger did much of the original work describing and categorizing the Magellan in-flight problem/failure reports. The use of many of his anomaly descriptions in this report is greatly appreciated. The author also wishes to thank Anthony Spear, Magellan Project Manager for his valuable comments, suggestions, and contacts for obtaining future information. The author especially wishes to thank John Slonski, Magellan Spacecraft Systems Engineer, for his careful review of the document and many useful suggestions for improvement. Thanks are also due to John Koch for many helpful suggestions on improving this report. The author also expresses his thanks to the JPL and Lockheed Martin personnel who are cited in the Bibliography for their comments and insights on many of the in-flight anomalies.

Correlation of the Magellan Flight PFR History with Ground-Test Results

SUMMARY

This report summarizes an investigation of in-flight anomalies that occurred on the Magellan Program and their relationship to environmental testing. An attempt is made to characterize or group the in-flight anomalies in order to find a common thread of lessons learned from Magellan. This document describes a method for evaluating in-flight anomaly data and extending the analysis to pre-launch test data, to identify test practices that are effective in precluding fight problems. By the same token, it is noted when product assurance elements other than testing had a greater influence on the occurrence of flight anomalies.

I. INTRODUCTION

Background

The Flight Anomaly Characterization RTOP has developed a method for characterizing and analyzing in-flight anomaly data to better identify where improvements in product assurance techniques and programs would benefit future flight programs. This method was based on the Problem/Failure Reports (PFRs) descriptions of the anomaly and the corrective measures taken. The study described here augments the flight PFR data with information from PFRs generated during the pre-launch ground-test program. Only anomalies occurring on the Magellan Program are considered here. A single program was selected because of the expanded scope of the study. It was felt that when all of the anomalies considered were from the same program, it would be easier to correlate the flight results with the test results.

Correlation of the flight anomaly history with the ground-test history provides insight into the reasons for disposition of the flight hardware when problems were identified during the development cycle. The hardware disposition, including the risk accepted after problems were revealed by testing and resulting analysis, can be critiqued. Recommendations made using only the flight data reports are likely to be based on the assumption that the appropriate test was not performed, or that an incorrect disposition of a test problem occurred. Investigation of the test history provides insight into other aspects of hardware development such as risk assumption. On the other hand, there are test related issues that were illuminated by working back through the pre-launch phase after considering flight performance. All issues were viewed in the context of the test program and why it was not 100 percent effective.

The Magellan flight program was selected because it was a recent program completing a successful mission, but with significant flight problems that provided an interesting set of data to deal with.

Objectives

The main objective of this study was to investigate the individual Magellan Flight PFRs to try to correlate them with the pre-launch developmental, functional, and environmental test history. Each of the PFRs can be considered as a missed opportunity somewhere in the hardware development cycle. The question of whether the test-program implementation, itself, was responsible for failing to uncover the potential problem, or some other phase of product assurance and reliability engineering was responsible, is the key issue considered here. A secondary issue is whether the potential problem should have been uncovered by the system-level testing or in prior testing. If this can be done, there will be a trail for potential discrimination among problems that are lower-level hardware-developmental in nature, i.e. parts, workmanship, or assembly-level or subsystem-level design issues, and problems resulting from system-level interface issues.

In those cases where testing is identified as the weak link, the next question is why the potential problem escaped identification. Possible reasons include:

a) the wrong diagnosis of a hardware problem;
b) the wrong corrective action, and in conjunction with this, failure to retest to verify the corrective action;
c) inadvertent failure to recognize the existence of a problem;
d) failure to supply the required rigor in the testing;
e) and waiving test requirements due to cost and schedule constraints.

There are flight problems where it is difficult to explain the inability of testing, on the basis of issues such as test level and duration per se, to uncover the potential problem prior to launch. There are other factors, and other functions in the product assurance cycle, that have a strong influence on the ability to prevent flight problems. These include both programmatic considerations and other product assurance functions such as analysis. These factors and functions will be discussed as they are identified. In addition, the various product assurance elements, including testing, can be considred to function together as a complete process to ensure hardware reliability.

Approach

The approach taken in this study was to use the flight problem/failures (P/ Fs), from the Magellan program, identified in a Fight Anomaly Characterization Task report (Reference 1) and attempt to trace the history of problems with the specific hardware associated with given flight P/ Fs. One source of data was the JPL problem/ failure report (PFR) data base kept by the Problem/Failure Operations (PFO) center. This also included copies of the problem/failure reports (also known as MARS) generated by the Martin Marietta Corp. (MMC). now the Lockheed Martin Corp. which will be referred to as LMC in this report. The MARS, generated by LMC during testing performed by them, are available as back-up information from the JPL PFO Center. Both electronically recorded data and hard copy back-up data and information, such as incidents/ surprise/ anomaly reports (precursors to PFRs) were used. In addition, personnel, both from JPL and the system contractorthat had worked on the Magellan Flight Program were contacted.

Two techniques were used to select problem/failure reports generated during functional and environmental testing. First, the actual back-up material available from the PFO Center for each flight PFR (including ISAs) contained references to test PFRs containing descriptions of test problems. The PFO Center database was also scanned to select PFRs that appeared to pertain to problems with assemblies or subsystems, that were similar to the flight problems. In addition to PFR and PFR-related data, other pertinent JPL documents and internal IOMs were used to obtain information.

II. DISCUSSION

There are a number of reliability engineering assessment tools available for screening spacecraft (S/ C) hardware problems before launch. Developmental, functional, and environmental testing and inspections at the assembly, subsystem and system level are important tools for uncovering potential flight problems. Most assembly and subsystem reliability-engineering activities and all system level activities are performed either by JPL or a system prime contractor. Other testing tools such as burn-in, life testing, and accelerated testing are applied mostly on a part or component basis, and frequently by the vendors and subcontractors involved.

A vendor typically has more than one customer for parts and components; thus, the vendor is relied upon to perform tests on his product. The vendor may not always do screening and life or qualification testing in a thorough enough way to prevent delivery of parts and components with potential problems. Other reliability assessment tools, besides testing, include design reviews and analyses, and failure modes and effects analyses, fault identification, complex circuit analyses, etc.

Twenty-six in-flight PFRs were investigated. Each one was assessed in terms of whether developmental or environmental testing at an assembly or higher level was deficient in terms of allowing a potential problem to escape detection. Also considered were specific issues such as screening and life testing; materials properties testing; requalification of inherited equipment; retesting and requalification after rework; and more thorough inspection of all similar hardware after a test failure.

III. PROBLEM SUMMARY

The following list of in-flight problem failures (P/Fs), by PFR number, provides descriptions, causes, fixes, environmental and other issues related to the PFR, as well as, a discussion of the probable reasons why the P/F was not prevented from occurring in-flight by the product assurance measures in place prior to launch. The combined impact of the PFRs on the mission and the reasons they occurred are used to derive the conclusions and recommendations presented at the end of the report.

PFR 52222: The Magellan Star Tracker (or Scanner) Unit (STU) provides attitude updates to the ACAS by acquiring two, known location, reference stars within its field of view (FOVEA). Beginning with the first star calibrations (starcals) attempted after launch, the STU began to generate unsuccessful starcals, hampering the spacecraft attitude update process. Review of star scan data suggested that the scanner signal was being prompted by a stimulus other than the target stars. This interference caused the starcal sequence to reject the correct star data, resulting in no attitude update; at other times, the false data was accepted, producing an incorrect attitude update. Starcal No. 1 rejected the first reference star and missed the second star, No. 2 misread the magnitudes of both stars, and No. 3 found both stars but rejected the first.

Analysis centered on the Starcal No. 3 failure. For each starcal in a series, the STU alternates between a forward scan (Swath 1) and a reverse scan (Swath 2). No. 3 was performed on the reverse scan sweep. Real time memory readout (MRO) analysis showed that the first star registered was the correct first target star (Gamma Crux), but that it was rejected because of a possible STU idiosyncrasy. It was postulated that because the reverse scan (Swath 2) caused the STU FOVEA to sweep across the Milky Way before reaching the target star, the background voltage buildup from first viewing the dense star field caused a misreading of Gamma Crux.

This explanation was later rejected when data received over the following 25 days showed that the next thirteen Swath 2 starcals were successful (with the exception of a single, unrelated failure). With the subsequent thirteen Swath 1 scans also successful, a random hardware failure within the STU also appeared unlikely. Further analysis led to a second preliminary conclusion that a stray particle or object passed through the STU FOVEA, possibly as a result of launch or near-earth activities, such as inertial upper stage (IUS) separation. This problem afflicted Magellan for over a year.

The STU demonstrated sensitivity to activity around the unit when it was turned on before launch-but the exact cause was never determined. There appeared to be a statistical correlation between activity ongoing near the hardware when it was powered-on, but no causal definition or relation to the testing per se. All analyses and testing indicated that the STU was operating as designed. No bad or degrading parts were found. Project personnel also felt that the expected-more-benign space environment would result in fewer problems than during development, and that any actual data loss would be acceptable. In addition to accepting the risk of any data losses when pointing errors caused loss of the target, they believed software changes were an option for mitigating any problems. There was no more test time to pinpoint the problem.

After contacting the vendor following the in-flight anomalies, JPL discovered that the star scanner was known to be sensitive to protons. Actually, the STU flight problems were related to two types of activity, where one was a sensitivity to energetic solar protons. Solar protons resulting from solar flare activity cause spurious STU interrupts. The effect varied from single-proton hits that simulated a star crossing the detector slit to a pulse of protons that looked like noise. The effects of the latter activity could be so great that star scan performance decreased to near zero for several days following an event. In effect, the sensor used on the star tracker was also a good proton detector.

The star scanner was designed for use in low earth orbit where it was shielded from solar protons. The specifications for the Magellan STU were written in terms of surviving exposure to radiation, which it did without a problem. However, they were not written in terms of sensitivity to proton radiation at the time of occurrence of an event such as a solar flare.

The second type of activity causing STU problems was related to reflection of sunlight off particles passing the field of view of the STU. In particular, the post-launch problem may have been related to shedding of Astroquartz®-fiber-insulation particles by the thermal blankets. Some shedding of the thermal blankets was observed after the system-level vibration test. The mechanism for this particle generation was not expected to pose a significant problem after launch and actually does not explain the continued problems in flight. However, the mechanism that generated the particles in flight was not fully understood; theories range from thermally induced mechanical stresses to solar wind particles which induced splitting of fiber particles causing them to break loose. The mechanism that expelled the blanket particles from the surface was believed to be the slight change in surface charge as areas of the vehicle transitioned from dark to light; the effect did not exist on the shadowed surface and was most noticeable during transitions. Use of these thermal blankets was forced by S/C thermal design decisions made on the basis of other factors-see discussion for PFR 52228.

Although, the noise from this activity was significantly greater in density and magnitude than the noise from solar protons, the impact of it could be controlled by imposition of constraints on dark-to light transitions. As state above, the effect became most noticeable during transitions from shadow into the sun. Occasionally this strategy did not work.
The sensitivity of the star scanner was lower than expected. This resulted from a mismatch between the S/C turn rate and the assumptions used to generate star catalogue provided by the vendor. LMC implemented a continuing software fix to mitigate problems, implementing a set of "foreground" and "background" software filters for the AACS star recognition process. The first filter, the "foreground" filter, was devised, two months after launch when problems first surfaced, as a filter on accepted scanner voltage. This filter was a magnitude check and was a trade-off between setting broad magnitude bounds accepting many spurious interrupts and narrow bounds blocking out real interrupts with voltage noise from proton impacts. The S/W patch also tightened up the period of time that the scanner interrupt was enabled. Later in flight a "background" filter, which screened out crossings that would result in an unusually large attitude updates was implemented. It picked out the best two of a maximum of eight crossings that gave the smallest resulting update. The motivation for this latter filter was to reduce the number of scans that were "thrown out". This led to the acceptance of some scans that should have been rejected; however the success rate was still over 95% for the scanner.

As stated above, a software change was a solution that was considered an option, if required, prior to launch. Although the type B activity (the Astroquartz® particles coming off the outside of thermal blankets) created too many interrupts to filter, a work-around was the proper selection of stars. Although the logic screening filtered out most of the spurious proton inputs, occasional problems recurred during solar flares-the mission took place during a peak in solar flare activity. The mission impact was considered minor, causing some swaths of missing data and some mistakes in correlation of the radiometric data with the radar data. The use of these fixes implies that developing work-arounds when problem/failure risk is assumed will help mitigate that risk.

One of the two fundamental causes of the problems with the STU, although not the worst (blanket shedding was worse), was the use of an inherited flight component not qualified for an environment (solar protons) different from its previous use. The same STU is being used on another flight mission, currently in development, with differences from the Magellan Mission. These differences include a trajectory away from the Earth with lower proton fluences from solar flares. Also, the spin rate of the spacecraft, in the current mission, allows more time in between star-location updates, so that there is more time to wait out a solar flare. This is providing, of course, that the mission is not right in the middle of a critical maneuver. This latter risk has to be accepted.

PFR 52223: During Magellan cruise, the motor current for gyro B-2 was seen to jump initially from 115 ma to 150 ma, and finally reaching a level of greater than 360 ma. There were accompanying jumps in temperature. Excessive gyro drift subsequently prevented DSN from locking onto the High Gain Antenna x-band for tape recorder playback. This was followed by further variations in current, temperature, and drift performance which led analysts to attribute the problem to a chattering bearing retainer in the gyro synchronous motor.

Spacecraft attitude control was then transferred to the alternate Attitude Reference Unit (ARU), which has been performing nominally. Gyro B-2 was eventually powered off due to the extremely high current levels (>360 ma), and the gyro vendor views it as a failed gyro. Diagnosis of the problem centered on an increase in the gyro motor torque caused by contamination or lack of bearing lubrication.

This in-flight gyro-bearing problem occurred because the vendor was not able to detect a flaw in pre-delivery cleaning of the gyro and notify JPL until after launch. The gyro lubricant became contaminated with solvent causing lubricant failure. The vendor believed that the Magellan gyros had demonstrated that they were not effected by the problem by having enough running time without failure. By this time, the problem with the fixture was corrected on new gyros that were being built. The problem resulted from a manufacturing-process change which allowed cleaning solvent, used during assembly, to leak around a fixture which did not always seal around an 0-ring. It should be noted that this type of bearing problem is a lifetime issue, and impossible to detect in non-accelerated environmental testing.

Although, gyro problems (overheating and erratic current) occurred during system solar thermal vacuum (STV) testing, it is not believed that they were related to the in-flight problem. The STV test-induced erratic behavior was of two types: 1) one resulted in excessive gyro temperatures; and 2) the second type resulted in excessive motor currents. In reality, the two problems were related and probably due to retainer-ring vibration, a problem , that if recognized by the vendor, was not revealed prior to launch. The temperature problem was treated as a heat conduction problem. At the time JPL was concerned that the excessive temperatures might be potentially life limiting. Post launch, it was finally determined that the erratic in-flight behavior was due to solvent contamination of the lubricant during assembly, and unrelated to the problems occurring during STV testing.

Therefore, the problem arose when a vendor made manufacturing-process changes, and did not thoroughly verify that the changes would not affect the functioning of the gyro. However, not all of the gyros built this way failed (prior to the vendor determining the cause of the problem), and this type of gyro is still in use.

PFR 52224: The command data subsystem (CDS) suffered a read parity error resulting from a bad memory cell. Fault protection software initiated S/C safing from Processor B, after switching from Processor A which went into a suspended state. The cause of the parity error was attributed to a gate oxide deficiency caused charge leakage allowing the memory cell to change the bit value. The culprit device was a Harris 6504 RAM. A software patch was used to change commands so that failed bit would not adversely affect the command sequence.

Although TCC244 RAMs experienced gate oxide failures during memory burn-in at Sandia, blocks of TCC244s were replaced with Harris 6504 RAMs in the CDS because of a eutectic die problem affecting the attachment of the wafer to the carrier. It was believed that mechanical failure could occur leading to fracture. The Harris chips which were processed differently and thought to be better. Although gate-oxide problems had occurred during bench tests of the TCC244 chips, and later with Harris 6504 chips, the risk of a significant mission impact due to gate oxide failure in a Harris 6504 chip was accepted to be minimal due to an expected low failure rate. Also impacting the decision to leave things as is, was the fully redundant memory.

This in-flight problem demonstrates one result of the classic trade-off between replacement of a part with demonstrated problems with one thought to be more tolerant to problems (although not problem-free). The Harris chips were not burned in or qualified as extensively as the TCC244 chips. The replacement part was also newer and consumed less power. The acceptance of the risk proved to be justified as this P/F had no mission impact.

PFR 52225: The DMS-A track 2 playback data showed corruption occurring in the frame-sync code and the data portion of the frame. The anomaly appeared on tract 4 next, and finally all four tracks were bad. This was also known as the "slip-a-bit, flip-a-bit" problem. The most plausible explanation for the DMS-A flight problems was low record signal strength to the tape, which was probably caused by an electronic component failure or degradation in the record electronics. The components considered were either of two specific capacitors which could have degraded via capacitive shorting. In-flight testing revealed that attempts to write over the tape failed to remove the old data. The reason accepted for the failure to eradicate the old data was that the voltage drop caused by a degraded capacitor was great enough to prevent a record relay from closing; thus producing no bias or data, allowing the old data to remain undisturbed. This failure mode has since been disputed and all that can apparently be concluded is that the problem was localized to the record electronics/head. In actuality, it could have been either a mechanical or electrical problem.

There were similar problems with this tape recorder in other satellite flights (the GSFC Compton Observatory (formerly the GRO satellite), Geosat and Hubble satellites), that may not have shown up in test. In the case of Magellan, the bit-flip error rate during pre-launch testing was acceptable. Therefore, the problems did not show up during pre-launch testing and check-out, and only worsened with and/or use during flight.

PFR 52226: Shortly after launch, telemetry indicated that Rocket Engine Module (REM) temperatures were higher than predicted. In a tail-towards-sun attitude, REM temperatures ranging from 39°C to 51°C were detected, instead of the 15 to 25°C expected. The direct operational impact of this anomaly was to impose attitude constraints on a certain portion of the cruise phase of the mission, as well as during mapping of the Venusian surface. Mission planners had to maintain spacecraft attitudes which would preclude pointing the rocket engine nozzles towards the sun. This precluded mapping of higher planetary latitudes during the first extended mapping mission, and it resulted in some loss of extended mission data.

Since an adequate subsystem thermal design appears not to have been done, the Rocket Engine Module (REM) problem was probably due to restricting the job of doing the thermal design on this subsystem to the REM vendor. Although, the system-level STV test was constrained by the S/C gimbal, an over-heating problem did appear during STV testing. The vendor implemented a corrective action (a sleeve around the 100 Lbf catalytic bed regions), but the schedule did not allow for test verification. Entrance of simulated solar insolation into the interior of the rocket nozzle in a tail-towards-sun spacecraft orientation, the principal over-heating mode of solar radiation going down the nozzle, could not be simulated with the existing test fixture. In addition, a low emittance nozzle coating impeded heat transfer. The test fixture was inherited from another program and could not accommodate an angle of incidence which would test this heating mode in the STV test. This was a case where a separate STV test at the assembly level allowing all of the various flight orientations would have provided a lot of useful data for evaluating the problem. A developmental model for use in such testing would have been useful.

The REM problems emphasize the importance of preventing future problems by performing a rigorous thermal analysis of the subsystem. In effect, the system contractor did not verify the thermal design to the extent required. And as stated in the "Lessons Leaned" Section in Reference 2, systems thermal vacuum testing cannot be substituted for component thermal design qualification or flight acceptance testing. This situation led to an expensive set of post-launch tests to verify operation of the REM at the higher temperatures.

PFR 52227: One month after the REM over-temperature problem was reported (PFR 52226), the temperature of the Solid Rocket Motor (SRM) forward flange area was also found to be higher than expected. Thermal analysis indicated that the effective solar absorptance of the separation flanges was underestimated. Model correlation showed that the absorptance was approximately 0.53 instead of 0.36. This additional heat transfer could be explained by an increase in the solar absorptance of the separation flanges above the value assumed by the Magellan TRASYS model. Subsequent tests showed that the resulting higher temperatures were within the acceptable range, and that no operational solutions such as attitude adjustments were necessary. Hence, this temperature limit violation had no impact on the primary nor the extended mission.

Because of the inability of the simulated solar radiation to reach the affected surfaces in the STV test due to limitations of the test configuration, the STV test was unable to detect this problem prior to launch. Therefore, it is important to develop the appropriate test fixtures and controls to test an assembly or subsystem in the expected flight configurations. In addition there was a lack of proper inspection of optical surfaces to ensure that they had the required optical properties.

PFR 52228: Beginning early in the mission, and continued through VOI and mapping, sensor readings in the attitude-control processor bay area, also known as the On-Board Computer (OBC), and the Command Data Subsystem (CDS) bay indicated higher than expected temperatures throughout the spacecraft bus. JPL attention focused on the 1 " x 1/2 " Optical Solar Reflector (OSR) tiles which control solar absorption by the Magellan spacecraft bus. Degraded reflectance of these tiles appeared to be the best explanation of the problem, since no defects could be found with the thermal model nor the solar thermal vacuum test data. Electronics bay temperatures were greater than predicted from early cruise, through Venus orbit insertion and mapping. Since the new OSR absorptance curve could affect flight operations, including limiting data collection during certain Venus mapping cycles, the mission impact was rated as "Potential for Major Impact."

The most likely cause for degradation of the OSRs is contamination of the reflective surface by spacecraft materials which affect heat absorptance. Major contributors to the contamination were the RTV adhesive used to attach the OSRs to the spacecraft and the structural adhesive used in the solar panel honeycomb. Organic products from this glue may have out-gassed when exposed to vacuum, and condensed on the colder surfaces (as compared to the solar panels which were warmer). Studies showed that the intense exposure to ultraviolet light may have caused them to polymerize on the OSR tile surface and reduced its reflectance causing an unfavorable absorptance (solar rad.) to emission (IR) ratio. Although Magellan thermal-vacuum tests did not reveal this problem, longer duration high temperature tests conducted by the TOPEX project on qualification panels using a similar honeycomb material detected significant venting from the edges of the panels. During aerobraking in the upper atmosphere of Venus, atomic oxygen apparently cleaned the OSR tile surface, lending credence to the "increased absorptance due to polymerization" hypothesis, rather than irreversible optical-surface degradation.

The problem surfaced early because of the need to reorient the REM; however, it would have occurred later in the mission in any case. Such high temperatures can exert a life-limiting effect on electronics. The thermal design of the S/C was marginal due to using fixed high-gain/ RADAR and altimeter antennas (where the altimeter antenna is also fixed relative to a particular point on the other antenna). This design required antenna pointing to be done through constrained S/C orientations so that all surfaces were exposed to the sun at Venus. Use of the fixed antennas led to a difficult thermal control problem.

Additionally, the fact that thermal control was compromised by contaminated optical surfaces and adhesives suffering chemical breakdown under UV exposure is indicative of the need for better materials control policies. These would include debugging before launch under flight-like environments, and a long (~50-100 hrs.) vacuum bakeout for removal of volatiles from the adhesives.

The "Lessons Learned" Section of Reference 2 states that "Full-scale thermal vacuum testing is mandatory for thermal design verification for establishing the interrelationships between assemblies and forfinalizing the thermal control configuration for the spacecraft [...]. The system-level verification can be supplemented through integrated thermal analyses and a sound test program involving components, materials, and partial configuration tests". All of this points to a lack of appropriate supervision of system-level thermal control. This lessons-learned statement also applies to the two PFRs discussed just prior to this one (PFRs 52226 & 52227).

A significant impact of attempting to control the peak temperature experienced in the electronic bays was an increase in the thermal cycling of the electronics. These control measures employed several attitude sequence scenarios developed to shade the electronics bays a portion of each 3.1 hour orbit. The bays in question housed the CDS and the OBC. The most utilized was the "two hide" concept which cooled the S/C by positioning the bulk of the S/C in the shadow of the high gain antenna twice each orbit. The objective was to keep the electronics below "flight acceptance" test levels of 50-55°C. (compared to desired pre-mission predictions of 30°C). The thermal cycling contributed to stressing of electronic packages including solder joint fatigue. An attempt was made to keep the depth of the thermal cycles to less than 10°C to minimize electronic solder joint fatigue that could cause premature failure. In addition, the performance of exciter B (see PFR 52243 below) tended to improve with increasing temperature and decreasing thermal cycling. Of course, the use of these strategies caused a reduction in mapping and playback time.

PFR 52229: This anomaly involved a change in the radar frame format that was not expected. The change in the format was due to a CDS/RADAR command timing-synchronization idiosyncrasy which was fixed by careful command timing. The exact command timing was never tested during ground tests; this might been avoided if all command sequences, along with their critical timing characteristics, had been tested in functional tests. Also useful, might have been a detailed review of intersystem interfaces. These kind of interface problems are amenable to system-level testing only. The problem was found and fixed during the cruise phase tests and thereafter, and did not have any impact on the mission.

PFR 52230: Following release of the two solar panels during near-Earth launch phase, Magellan telemetry provided no initial indication that the panels were latched. At the same time the Shuttle astronauts observed the panels to be deployed. The microswitch on each panel must close to provide a latch indication. The panels were then rotated into a position where they received a positive reaction force during the launch vehicle upper-stage burn, providing a higher probability of deployment. A solar panel latch indication was received a few seconds after engine ignition, so no further action was required.

PFR 52231: This PFR consisted of the first five separate incidences of many TWTA shutoffs, starting 4.6 months after launch, that were eventually mitigated by operational work-arounds. Fortuitously, they did not occur during a critical time such as Venus Orbit Insertion (VOI).

Spacecraft fault protection was eventually modified to make it's initial response to a loss of TWTA output be a restart of the same TWTA, rather than a swap to the backup TWTA.

There were four such TWTA incidents before launch during ground testing. No long-term corrective action was recommended. The TWTAs on Magellan were inherited from the International Solar Polar Mission (ISPM, later Ulysses) and had heritage with the GSFC TIROS/NOAA and Air Force DMSP meteorological satellite programs and were not requalified for the Magellan mission. The fact that TWTA shut-offs occurred in the flight acceptance tests indicates that an inherited design characteristic was at fault. Although, the shut-off were accepted as an operational idiosyncrasy and caused only a small amount of data loss, they occurred more than predicted and were of some concern. A detailed review of inherited designs and possible requalification of inherited hardware for Magellan requirements that may have been different than prior programs may have been beneficial to the program.

PFR 52232: A 0.5 amp deviation in the +X solar panel output, as compared to the -X panel output, was detected 4.4 months after launch of Magellan. The timing of the power loss was coincidental with a penumbral spacecraft alignment placing the altimeter antenna (ALTA) structure in front of the solar panel in line with the sun. Data suggested that the ALTA was casting a shadow onto the lower portion of the +X panel, reducing power generation. Review of pre-launch photos and drawings showed such an overlap. Due to an adequate power margin, the loss of 0.5 amps when the ALTA was in front of the solar panel was viewed as minor and as having no mission impact. The cause of the anomaly was insufficient analysis of structural interference (structural design).

PFR 52233: This anomaly involved the improper functioning of the range/rate time mechanism, in effect disabling the rate/range and gyro variance detections. A software fix was used to set the rate switching delay timer preset to 0. It appears to be a software issue, rather than an assembly-level or system-level testing issue. This had no impact on the mission.

PFR 52234: Magellan telemetry provided an intermittent indication that two channels on gyro B2 were producing gyro counts at full scale. After gyro power became reset during the autonomous response to the first heartbeat loss anomaly (see PFR 52236 below), the B2 outputs were observed to be nominal and consistent with readings from other gyros. Attempts to reproduce this failure mode were unsuccessful, and the cause is unknown. The only corrective action implemented was to reassign the B2 channels to backup use.

PFR 52235: Within 7 seconds after solid rocket motor separation, the B-side of the Magellan command data subsystem (CDS-B) received erroneous alert codes from AACS-B. This anomaly was isolated to memory B in the on-board computer, where memory bit 4 was found to have stuck high, causing the read/ write operations to mis-address a block of memory used to receive CDS commands. This addressing failure affected at least 2K of RAM. Memory B was marked off-line to inhibit read/writes to memory B, preventing the AACS from operating in memory B RAM (preventing an inadvertent command to be accepted by the AACS). JPL was able to match these symptoms using a failure model in which a latch up failure occurred to a TCC244 chip.

A voltage transient through the spacecraft chassis is the suspected cause of the memory failure. JPL determined through ground tests that by firing one or more NSI (NASA standard initiator) devices, a plasma path to the case could conduct enough chassis current to cause the memory failure. Chassis current can be generated when a post-fire plasma in the NSI allow the NSI conductors to short to the case during their firing. Eight NSIs were used simultaneously during

SRM separation. The AACS memory board is physically located 1/4-inch above the ground plane, and a voltage transient of only one volt is sufficient to cause the AACS memory B failure. The results of NSI ground test firings led to the conclusion that a memory failure could result from this noise-induced environment, and a TCC244 latch up model prediction of eventual "self-healing" corresponded to observations.

Magellan was equipped with two redundant AACSs, including two 32K memories and two processors-all cross-strapped to be interchangeable. When the memory loss occurred, memory B was serving as a backup, performing the same functions as memory A but not controlling the spacecraft. Although some areas of memory could not be examined from the ground, it appears likely that 4-8K of memory B became unreliable and, in effect, unavailable for use. If this loss of available memory space had occurred <u>permanently</u> in memory A and the spacecraft had had no redundant B-side, the mission would have ended at SRM separation. Assuming an <u>intermittent</u> condition occurring later in a single-string mission, ground controllers might have been able to program around the glitch. Their success would have depended on which particular code was in the affected memory space at the time of the failure and whether the programmers had time to insert a fix before the mission entered a critical phase.

The cause of this memory problem was not commonly recognized at the time of occurrence. This is basically a circuit design issue, that was not easily recognized, involving pyrotechnically-initiated electromagnetic interference that ultimately resulted in a latch-up of the electronics. It is recommended that it be treated in the future as a circuit-design issue requiring a thorough analysis. There is further discussion in the recommendations section.

PFR 52236: This problem occurred when the command data system (CDS) detected a lack of Attitude Articulation Control System (AACS) heartbeat, and began the heartbeat loss response execution, putting the S/C into a safing mode. The lack of a heart beat is attributed to a runaway program execution (RPE). The RPE is now considered to have resulted from a bug in the flight-control software. This problem was exacerbated by memory problems in the backup computer caused by latchup tied to a pyro-firing induced by ground currents (see problem discussed above for PFR 52235).

The direct immediate cause of the RPE was a timing idiosyncrasy in the AACS flight software that caused the AACS to enter an infinite loop. Factors which possibly exacerbated the problem were the recurrence of a non-repeatable pre-launch back-up computer problem, plus use of the bad memory of the backup computer, i.e. caused by the latchup. The non-repeatable pre-launch back-up computer problem caused it to randomly walk through memory. This combination of these three factors, the RPE, the back-up computer with an intermittent, and the faulty memory due to a latchup, all came close to causing mission failure. Eventually, the flight fault-protection software was modified to enable a rapid restart of the on-board computer by giving it a hardware reset without entry into safing or cancellation of the sequences. A software patch was also used to temporarily circumvent the bad memory which eventually annealed and healed itself.

PFR 52237: This was another example of the tape recorder problems discussed in connection with PFR 52225 above with problems leading to data corruption.

PFR 52238: The writing to a protected memory can probably be attributed to a runaway program execution (RPE). See the discussion for PFRs 52236 above, and 52241 for the causes of and the responses to the PFR.

PFR 52239: A gyro swap occurred during Venus orbit insertion and separation of the SRM due to a coincidental convergence of separation dynamics with stepper-motor induced rate changes and a tight variance threshold value. Since it had no impact on the mission, no action was required. This involved flight hardware and software interactions, not amenable to pre-launch testing.

PFR 52240: During Venus encounter, temperatures in the Magellan electronics bays exceeded Venus orbit insertion (VOI) predictions and bay acceptance limits. This problem was fully discussed in the comments on PFR 52228.

PFR 52241: This was another occurrence where the CDS detected a lack of AACS heartbeat, and began the heartbeat loss response execution. Unlike the problem discussed for PFR 52236 above, this problem only involved the RPE. The backup computer was not used and there was no memory latchup. Therefore, the mission was not put in jeopardy as before.

PFR 52242. The Magellan solar panels began jittering during mapping passes, causing the spacecraft to oscillate. Analysts noted a growing divergence since the beginning of mapping operations between the solar array drive motor (SADM) commanded position and the potentiometer reading of actual position. If this divergence had been allowed to continue, flight software would eventually have signaled a SADM Control Loss fault indication. JPL attributed the divergence to torque applied to the drive mechanism by the repeated changes in the direction of panel movement during jitter.

The problem was exacerbated by the excitation of a structural resonance in the solar panels, causing severe vibration. The jittering effect itself, however, was caused by a deficiency in the flight software algorithm used to calculate the desired panel position for oblique sun incidence angles. This problem was corrected with a patch to the articulation control flight software, eliminating the jitter by switching the solar array control to an open-loop mode (not using sun sensors while celestial geometry was unfavorable).

PFR 52243: This PFR occurred 22.5 months after launch and resulted in significant downlink degradation due to the presence of a sweeper spur caused by a failure in the last stage of the X-band downlink exciter B; capability was down to 43%. The failure seriously degraded the transmission capability by preventing the subcarriers from operating. Transmission was switched to exciter A at this point, and therefore, the only consequence was a significant loss of redundancy.

Although the cause of the failure was unknown, the most likely cause was a damaged chip capacitor, a CDR01 type, in the last stage driver circuit. It was postulated that value of the capacitance changed due to sensitivity to thermal stresses; such capacitors often crack when over-stressed thermally. The thermal stress involved here was a large number of thermal cycles after the stress of interplanetary vacuum. Prior to launch, the capacitor underwent a number of soldering operations leading to the possibility of damage during the assembly and alignment of the module.

The reason for the rework was a warning from the radio manufacturer that some of the internal ground straps did not have sufficient stress relief and had a high probability of breaking. Upon disassembly, broken straps which would have led to potential flight failures, were found.

The capacitor did not exhibit degraded performance (lower capacitance) until after the S/C was subjected to a relatively large number of flight thermal cycles. The transponder was qualified to have a lifetime exceeding that experienced during flight; the chip-capacitor lifetime was shortened by the thermal stresses during the above-mentioned rework. A better requalification plan for, or inspection of, components after significant rework n-Light have mitigated this problem. It would also be beneficial to minimize the amount of rework on flight components.

PFR 52244: Five Magellan temperature-sensors failures were documented on a single PFR (PFR No. 52244). These temperature sensors failed in various locations at various times starting approximately 1-1/2 years after launch:

Part No. ISA No.	Anomaly Date	Location	Symptom
E-00149053	90-295	Multi Layer Insulation T2	Erratic Performance (SADM)
E-01009054	90-2%	HGA Reflector Sensor	Erratic Performance (SADM)
E-MS8766	90-336	Multi Layer Insulation T3	Failed High (1255dN)
E-01068761	91-M	Medium Gain Antenna	Failed High (255dN)
E-01849280	91-218	Shunt Radiator T1	Erratic Performance

Three of the sensors began giving erratic readings (two of which were correlated with solar array drive vibration) and the other two read full scale, indicating an open circuit. Sensor E-0100, listed above as exhibiting erratic performance, also failed later as an apparent open circuit. Prior to the failure of Medium Gain Antenna sensor E-0106 during orbit 1733, it was cycling approximately 155^0C per orbit and had been performing erratically. HGA reflector sensor E-0100 was also analyzed: it began to exhibit erratic performance 17 months into the mission after experiencing both extreme temperature cycling of $100\text{-}150^0$C and solar array drive motor (SADM) vibration. No details are available on the environment of the other three anomalies.

The thermal sensors in question were not measuring temperatures whose loss represented a critical data loss to the mission at the time of detachment. In addition, they were placed in locations where the risk of detachment was significant. The environment at the point of attachment, especially dynamic environments, make it difficult to sustain the attachment. Some thermal sensors were also mounted on the thermal blankets, a particularly difficult location to secure such a sensor. When thermal sensors are mounted internally to measure a critical function, several sensors are used to provide data which can then be cross-correlated. The chance of the internal ones falling off is much less than of the externally placed sensors. In the event that future missions will require critical temperature measurements in severe environments, better attachment techniques will have to be developed and qualified to these environments.

PFR 52245: The radar pre-regulator current sensitivity to a soft ground violation was attributed to a propellant line heater short. This had no impact on the mission and there were 3 remaining parallel heaters to make up for the short fall. A similar problem occurred in ground testing, caused by a clamp on the propellant line cutting into the heater lines causing a short. There is no explanation why the testing did not detect the potential for the second problem which apparently required the added exposure to the flight environment.

PFR 52247. This was a case of DMS motor fluctuations that were determined to be normal and an incident surprise anomaly was voided.

PFR 52621: This RFS PFR occurred 2.7 years after launch when the subcarriers on exciter A were lost. The S/C was switched over to exciter B which had only 43% (see PFR 52243) capability. At this point the primary objectives were met; further RADAR mapping was precluded, but, since the main carrier was still good, the gravity science measurements could be accommodated.

The failure was attributed to leaky glass passivation on an op amp that allowed trapped moisture to form an acidic compound that found its way to a resistor, etching it and causing it to fail. Although the passivation system was suspect and a failure occurred in ground testing, schedule problems and the lack of significant data on a new passivation process at the time restricted the use of devices with the new system. Also, rigorous testing of the older devices in the same lot in water produced no failures. Devices with the older passivation system are no longer used.

The loss of the transponder subcarriers was due to a part failure. However, the part was rigorously tested before launch with only one failure. The op amp in question had a passivation technology no longer used; however, schedule constraints at the time prevented replacement. This failure may have been caused by the power/thermal cycles the transmitter received during mapping (see discussion for PFR 52228 above). Also, the rework undergone by the transponder may have impacted the lifetime of the part.

IV. SUMMARY OF LESSONS LEARNED

System Contractor/Subcontractor/Vendor Interaction

The Rocket Engine Module (REM) overheating resulted from the system contractor failing to validate the thermal design of the vendor and not insisting on rigorous testing at the subsystem level under solar illumination. This is an example of what can happen when the system contractor assigns a systems-level integration function to a vendor or subcontractor and does not validate the results. This lack of attention early on, coupled with the limited STV test performed and lack of proper retest on the fix imposed, led to a very expensive post-launch program to verify REM operation at the temperatures being experienced. The STV test was limited in that the REM could not be tested at multiple attitudes with respect to the sun. Since the actual REM operating environment was tied to the fixed high-gain/ RADAR and altimeter antenna design, the necessity of thoroughly thinking through a design and its subsequent impact on hardware operation and testing cannot be overemphasized.

There are several examples where closer interaction with vendors and subcontractors might have prevented or alleviated problems. The drive by LMC to hold costs on a low-cost fixed-price contract may have impeded data flow and working relationships. The gyro, which had a bearing problem, was one example. Problems showed up in environmental testing. The cause of the problem was attributed to a retaining ring in which a materials change had been made. The actual cause of the bearing wear was due to a manufacturing-process change that resulted in tooling being used during the cleaning process that allowed solvent to leak into the bearing, contaminating the lubricant. The vendor had not qualified the gyros produced using the new tooling.

The tape recorder was another example of a component with problems in flight. Both the tape recorder and gyros may have benefited from better life testing.

Thermal Control Problems

The thermal control problem on Magellan was particularly difficult because of the necessity of exposing all surfaces of the spacecraft to the sun at Venus, which is the equivalent of two earth suns. Exposure of all surfaces to the sun was necessitated by the use of fixed high-gain /RADAR and altimeter antennas (i.e. pointed these antennas at Venus, exposing most S/C surfaces to the sun during most of the orbit). The use of a sun shade as the S/C orbited around Venus was prevented. This required S/C thermal-control strategies to reduce peak temperatures in some of the electronics bays that led to some undesirable thermal cycling and loss of science data. These latter effects of the thermal-control strategies were mitigated by demonstration, through margin testing, of the ability of the electronics to operate for periods of time near 50°C.

The Magellan S/C was forced to operate under severe thermal conditions (i.e., subjected to both higher temperatures and transients associated with orbital exposure variations). This condition was made more difficult because of the unanticipated changes in optical surface properties during flight, as well as inadequate characterization of some of the optical properties of exposed metal surfaces. The former condition involved contamination of the optical surface reflectors (OSRs), and the optical-surfaces coating by chemical compounds formed from UV breakdown of the adhesives holding the OSRs. The problem with metal surfaces resulted from a lack of adequate inspection criteria to verify that the desired optical surface properties (emittance value measurements) were achieved as specified by thermal control engineers. This was a workmanship/ processing problem that both LMC and JPL missed. There were also some constraints on the STV testing that impeded adequate characterization of expected thermal performance.

In summary, the thermal design , and its associated transient characteristics, was difficult to analyze and test, and this made it difficult to correlate ground-test results with the thermal model and materials properties.

Inherited Hardware

The TWTAs on Magellan, which were inherited from the ISPM program and had heritage with GSFC TIROS/NOAA and Air Force DMSP meteorological satellite programs, were not requalified for Magellan. The occurrence of the TWTA shut offs during the flight acceptance tests suggests that the acceptance of an inherited design without requalification may not have been a wise decision. That is, requalification of the TWTAs would have occurred earlier in the development cycle and thus may have provided more time for addressing the corrective actions for the problem. However, cost constraints may have prevented any decision that differed from those actually implemented. This would point to the necessity of a detailed review of inherited designs and the need for requalification of the hardware for different uses than originally intended, including different requirements and environments.

The Star Tracker Unit (STU) was inherited from the Inertial Upper Stage (IUS) launch vehicle. In that application, it was only subjected to particle radiation present within the earth's Van Allen radiation belts. It was also calibrated with a star catalogue appropriate to the turn rate of the IUS vehicle. Use of this hardware on the Magellan S/C subjected it to solar protons produced during solar flares. The STU had been qualified to survive protons, and to function after exposure but not during the proton flare event. This failure to qualify the STU to function during a solar flare led to spurious responses during proton flares. The improper star catalogue problem was less of a problem but required updating of the catalogue.

Impact of Faulty Flight Software

Although the main emphasis of this report is on hardware problems and hardware product assurance, the first occurrence of the loss of AACS heartbeat is an example of hardware problems interacting with a software problem, to create a situation much more serious than if the problems had not occurred coincidentally. The event was initiated by a software timing idiosyncrasy causing the flight software to enter an infinite loop; the S/C safing sequences combined with a computer with an unrepeatable, possible hardware, fault and a latchup in one of its memories almost caused the loss of the mission. The latchup was probably due to a pyro-firing. The initiating problem was a software problem, since the hardware anomalies, by themselves would not have led to the potentially catastrophic consequences. This is, therefore, an example of a potentially mission-threatening software problem, with possible hardware involvement, that must be identified prior to launch.

Pyro-induced Structure Currents

A problem not previously well understood surfaced during the Magellan flight. This was the occurrence of pyro-activation induced latchup. This resulted from the electromagnetic coupling generated by unintended transient-current ground loops inducing noise voltage transients in sensitive electronic circuits. The sequence of events leading to the transient ground-loop current is initiated by the heating of a bridge wire setting off an explosive charge. This can result in cracking of the encasing ceramic cup of the activation device, a squib, and generation of a plasma shorting the bridge wire to the chassis. The resulting transient current loop will induce a voltage in a near-by circuit loop (containing the electronics), oriented so that the magnetic field of the former has a component perpendicular to the plane of the latter.

Preventing this type of problem requires careful circuit design and physical placement of circuitry (See Reference 4). Two approaches can be used to prevent this problem. The first is to eliminate the structure current by isolating the return circuit-from ground by placing a high resistance (~5 KOhm) between ground and the power source. It should be noted that Magellan had such a resistor, but had a capacitor in parallel to deal with RF noise which could have defeated the use of the resistor. If isolating the pyro-activation circuit is not practical, then the effects of the structure current must be minimized. On the S/C, this can be achieved with placement and layout of the sensitive circuits and positioning of the pyro circuits. This should be followed up by a detailed analysis, and/or a test program. If testing is performed, care is required so that events do not occur that could prove harmful to electronic circuits.

In summary, the issue of pyro-induced structure currents is primarily a design layout problem and a ground-current isolation problem. It is prevented by adhering to good design practices in laying out the electronics, and isolation of high (pyro) currents. An engineering-development model should be tested by firing a simulated (non-explosive) dummy pyro device.

V. CONCLUSIONS AND RECOMMENDATIONS

There is one thread through the pre-launch hardware performance history as documented in Reference 3 and reinforced by the study being reported upon here. That is that cognizant engineers and S/C system managers at any mission-or-program-responsible agency, such as JPL, need to work very closely with the system contractor to assure that problems that occur during development are adequately addressed and rigorous corrective actions are implemented. In turn, the system contractor, needs to continuously interface with vendors and subcontractors to also stay on top of these problems. This follows from the lessons learned on Magellan in dealing with the problems on the star scanner unit, tape recorder, gyros, and REM, for example. Reference 3 makes the following recommendation: "For each work unit, a subsystem cognizant engineer should be assigned to follow the subsystem progress from requirements and design to delivery. This is especially important for one-of-a-kind builds. Subcontractors should follow this procedure also." Reference 3 cites the tape recorder and the transponder/exciter as cases where the chief design engineers did not follow through to development. Both subsystems had hardware failures. The structures/mechanisms and propulsion are cited as subsystems where the cognizant engineer did follow the subsystem throughout the process. "These subsystems had fewer and less serious hardware problems."

When assembly and subsystem testing and analysis are performed by vendors and subcontractors without the proper control, compromises may arise. This is further aggravated when design and manufacturing-process changes, however minor, on components such as the gyros, are made without the proper screening by the vendors. System-level testing is limited in its ability to screen out problems resulting from inadequate lower-level testing. For example, attempting to screen out problems in system and integration tests may cause conflicts with scheduling. There may also be a reluctance to undergo rework that may induce new problems. All of this tends to lessen the benefits of a good overall test program.

This report also addresses the programmatic lessons learned from the study of the Magellan records and flight performance. These important lessons learned relate to the design, development, fabrication and test processes. The following paragraphs provide conclusions and recommendations relative to these lessons learned.

<u>Do Problem Prevention Early</u>

Although it may seem like a truism, the Magellan experience reinforces the concept of doing things early in the hardware development cycle that will reduce problems in subsystem-and-S/C-level testing and flight. The occurrence of problems at the system-test level frequently results in insufficient time to retest, rework, and fix and requalify hardware with part substitutions and design and manufacturing-process changes. Examples of flight hardware where this was a factor are the star tracker, the gyros, the rocket engine module, and the TWTAs. Preventing this requires that JPL managers and cognizant engineers work closely with the system contractor and subcontractor to resolve problems early in the process.

<u>Inherited Hardware</u>

Programs like Magellan rely on inherited hardware design and components. Examples of problems with inherited hardware have been described above. The prime example was the problem of the star scanner which experienced serious problems in flight. The star scanner was never qualified to operate outside of the earth's Van Allen belts. It had been qualified to survive protons, and to function after exposure but not during the proton flare event. Other cases of problems with inherited hardware or hardware shared with other flight programs include the TWTAs, the gyros, and the tape recorders. The gyro problem resulted from the use of a gyro, whose manufacturing process was modified by the vendor and not thoroughly checked out.

Any design, processing or materials changes in components, such as occurred with the gyros on Magellan, should be scrutinized. Attention should also be paid to components where problems have arisen in testing or in other flight programs, such as was the case with the Magellan tape recorder and TWTAs. This will require working with vendors, and subcontractors to ensure that potential problems are screened out prior to finalization of the design process and fabrication. Design reviews, screening tests, burn-in tests, life tests, and later in the process, assembly and subsystem tests are all tools that are useful.

Cost and schedule pressures sometimes make it expedient to accept inherited hardware without requalification when changes are made to the hardware, or to the functional performance and environment in which it must work. This temptation must be resisted and vigilance exercised in screening parts, components and assemblies on inherited hardware.

In general, problems related to operating time, such as with the tape recorder or gyros may be uncovered by better life testing. In addition, those problems that do not surface until flight may be analysed to the point where thay are not considered to pose a significant risk to the current mission (because of redundancy or some work-around) and dropped. However, they may surface in a future mission and pose a significant risk under the conditions of that mission. Therefore, it would be beneficial to future missions to carry problems to conclusion and attempt to find causes for problems that could pose significant risk.

System-level Design

System-level design issues are another problem that require particular attention from a systems contractor. A significant problem experience by Magellan was the result of LMC delegating aspects of the system-level thermal design effort to a subcontractor. That was compounded by not doing the subsystem-level testing required to prevent problems from surfacing during system-level testing or in flight. In many cases this will require subsystem testing under flight environments. Although, it may be more costly to do the additional testing, valuable time and additional cost will be saved later, as well as the achievement of a reduction in flight problems.

All of the S/C maneuvering was not planned for prior to launch since, although Magellan was anticipated to be warmer than a S/C orbiting a planet further from the sun than earth, it was not expected to the extent that actually occurred during the mission. Lack of use of a shading strategy (sun shade) that did not involve transient thermal control leading to thermal cycling exacerbated the thermal problems on Magellan

Margin Testing

The Magellan experience proved the value of margin testing. Flight hardware was forced to operate at higher temperatures than predicted. System-electronics operation at approximately 50°C was demonstrated during STV. When optical properties degraded, extended operation near 50°C required fewer attitude maneuvers to shade the S/C in order to cool the electronics below 50°C.

Demonstrating that the REM could operate at the higher temperatures occurring in cruise was expensive, especially since it was done after launch. In order to do thermal characterization of a S/C it is essential to address the high temperature at the inner planets or the potentially cold temperatures at the outer planets.

Hardware/Software Interactions

In connection with the pyro-induced hardware problem discussed in the previous paragraphs, a software timing idiosyncrasy (bug in the S/W) created a mission threatening situation. This illustrates the need for careful pre-launch validation of software, and a full accounting for the effects of any software bugs coincident with hardware anomalies.

Finally, even if all of the correct design, test, and fabrication decisions are made, hardware failures may still occur late in the S/C development phase. At this point, trade-offs will have to be made between the expenditure of time and resources to fix problems, and the risk of flying hardware with potential problems. This will require new analysis on, or a review of the history of, the hardware in question. A prediction of flight performance will have to be made and a decision made on risk acceptance.

REFERENCES

1. Brown, A.F., “Development of a Method for Flight Anomaly Characterization”, JPL D11382, January 1994.

2. Martin Marietta Technologies, Inc., Flight Systems, Denver, CO, “Magellan Spacecraft Final Report”, JPL Contract 956700, MGN-MA-011, January, 1995.

3. Ledbetter, K.W., Wall, S.t)., “Magellan. Lessons Learned-Proceedings of the Magellan Lessons-Learned Workshop Held December 1991”, JPL D-9643, April 14,1992.

4. Leung, P., Nguyen, T., “Transient Structure Current Generated by the Activation of a Squib Device”, Proceedings of the 1995 IEEE Aerospace Applications Conference, Snowmass CO, February 4-11,1995, Volume 2, pp. 145-154.

BIBLIOGRAPHY

1. Oberhettinger, D., “Investigation of Thermal Sensor Failures Aboard Unmanned Spacecraft”, JPL D-11377, March, 1994.

2. Brown, A.F., “Analysis of Uplink/downlink Anomalies on Six JPL Spacecraft”, JPL D11383, September 1994.

3. Oberhettinger, D., “Investigation of Mechanical Anomalies Affecting Interplanetary Spacecraft”, JPL D-11951, September, 1994.

4. Oberhettinger, D., “Investigation of Environmentally-Induced Anomalies Aboard JPL Spacecraft”, JPL D-12546, March, 1995.

5. Nguyen, T, “MGN Memory B Failures: TCC244 Latchup Hypothesis”, JPL IOM 521190-374, September 21, 1990.

6. Gibbel, M., “MGN Spacecraft (S/C) Solar Thermal/Vacuum Margin Test (STVM Test) Preliminary Summary”, JPL IOM 5214-88-109, September 29, 1988.

7. Quinn, J.D., “Flight P/FRs and the Design Review Process”, JPL D-11381, January, 1994.

8. Clawson, J.F., “Lessons Learned from MGN Ground Thermal Tests and Flight Experience”, JPL IOM 5214-90-36, April 3, 1990.

9. Clawson, J.F., “Effects of Higher-Than-Anticipated Temperatures on the Equivalent Life of the Magellan Electronics”, JPL IOM 5214-90-99, November 6, 1990.

10. Clawson, J.F., “Magellan Lessons Learned”, Technical Presentation, October, 21, 1993.

11. Odetics, Anaheim, CA, “Ground Testing and Analysis of Magellan DMS A Data Degradation”, Report No. 9114026, Rev. NC, January 13, 1992.

12. JPL PFO Center Database, Electronic and Hardcopy Backup, of JPL Ground-Test and Flight Problem/ Failure Reports (PFRs), LMC MARs Failure Reports, and JPL Incident Surprise/Anomaly Reports (ISAs), March-April, 1995.

13. Gonzalez, Charles C., “Relationship of Test Program History to Flight History-Voyager and Magellan Radio Frequency Subsystem”, Significant Trend Report-TETA-TO-0024, August 26, 1994, Environmental Test Effectiveness Analysis Reports: For NASA Office of Safety and Mission Assurance (Code 0) JPL D-11295, REV. B, December, 1994.

14. Seale, E.H., “Magellan Star Scanner Experiences: What a Long Strange Trip It's Been...”, AAS 91-072, paper presented at the 14th Annual AAS (American Astronatutical Society) Guidance and Control Conference, February 2-6, 1991, at Keystone, CO.

15. “Magellan In-flight Solar Array Deployment Indication”, File No. 7-109, Submitted by B. Wagoner/D Sevilla, 11/5/90, JPL Lessons Learned File.

16. “Magellan Solar Array Drive Jitter”, File No. 7-110, Submitted by H.K. Bouvier/J. Blosiu, 9/3/91, JPL Lessons Learned File.

The following private communications provided information that was used in writing this report:

a. Private communication with Joseph Plamondon, JPL, regarding Magellan thermal problems, March 31,1995.

b. Private communication with Kenneth Starnes, Martin Marietta Corp., regarding Odetics tape recorder, April 4,1995.

c. Private communication with Julie Webster, JPL, regarding Odetics tape recorder, April 3,1995.

d. Private communication with Douglas Griffith, JPL, regarding Odetics tape recorder and also general lessons learned, April 3, 1995.

e. Private communication with Tien T. Nguyen, JPL, regarding various part problems and also pyro-induced structure currents, March 30,1995.

f. Private communication with James Clawson, JPL, regarding Magellan thermal problems, March 31,1995.

g. Private communication with Edward Liddy, and Russell Allen, JPL, regarding the gyro problems, April 4,1995.

h. Private communication with Paul Gordon, JPL, regarding the REM overhearing problem, April 5, 1995.
i. Private communication with Gary Parker, JPL, regarding the gyro problems, April 5,1995.

j. Private communication with Paul Gordon, JPL, regarding the heating-line clamp cutting into electrical insulation causing a short, April 17,1995.

k. Private communication with Joan Feynman, JPL, regarding the STU sensitivity to protons, Sept. 13, 1995.

1. Private communication with Tony Spear, JPL, regarding various P/ Fs and other contacts for information, Sept. 13, 1995.

m. Private communication with John Slonski, JPL, regarding of the STU anomalies, Sept. 14,1995.

n. Private communication with Frank McKinney, LMC, regarding various P/ Fs including the STU, gyros, tape recorders, heartbeat loss, and the OSR tile contamination, Sept. 15,1995.

o. Private communication with Jim Newman, LMC, regarding the OSR tile contamination, Sept. 15,1995.

p. Private communication with Eric Seale, LMC, regarding the STU anomalies, Sept. 18, 1995.

q. Private communication with Rick Kasuda, LMC, regarding the heartbeat-loss anomalies, Sept. 18,1995.

r. Private communication with John Slonski, JPL, regarding various anomalies and the conclusions drawn from them, October 13 & 17,1995.

s. Private communication with Greg Levanas, JPL, regarding anomalies with the CDS Harris 6504 RAMs and the tape recorder, January 5, 1996.